arXiv:1603.07181v2 [cs.IT] 19 Sep 2016

Iterative Scaling Algorithm for Channels

Paolo Perrone*[1] and Nihat Ay[1,2,3]

[1]Max Planck Institute for Mathematics in the Sciences, Inselstrasse 22, 04103 Leipzig, Germany
[2]Faculty of Mathematics and Computer Science, University of Leipzig, PF 100920, 04109 Leipzig, Germany
[3]Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

Abstract

Here we define a procedure for evaluating KL-projections (I- and rI-projections) of channels. These can be useful in the decomposition of mutual information between input and outputs, e.g. to quantify synergies and interactions of different orders, as well as information integration and other related measures of complexity.

The algorithm is a generalization of the standard iterative scaling algorithm, which we here extend from probability distributions to channels (also known as transition kernels).

Keywords: Markov Kernels, Hierarchy, I-Projections, Divergences, Interactions, Iterative Scaling, Information Geometry.

1 Introduction

Here we present an algorithm to compute projections of channels onto exponential families of fixed interactions.

The decomposition is geometrical, and it is based on the idea that, rather than joint distributions, the quantities we work with are channels, or conditionals (or Markov kernels, stochastic kernels, transition kernels, stochastic maps). Our algorithm can be considered a channel version of the iterative scaling of (joint) probability distributions, presented in [1].

Exponential and mixture families (of joints and of channels) have a duality property, shown in Section 2. By fixing some marginals, one determines a mixture family. By fixing (Boltzmann-type) interactions, one determines an exponential family. These two families intersect in a single point, which means

*Correspondence: perrone@mis.mpg.de

that (Theorem 1) *there exists a unique element which has the desired marginals and the desired interactions.*

As a consequence, Theorem 1 translates projections onto exponential families (which are generally hard to compute) to projections onto fixed-marginals mixture families (which can be approximated by an iterative procedure). Section 3 explains how this is done.

Projections onto exponential families are becoming more and more important in the definition of measures of statistical interaction, complexity, synergy, and related quantities. In particular, the algorithm can be used to compute decompositions of mutual information, as for example the ones defined in [3] and [4], and it was indeed used to compute all the numerical examples in [4]. Another application of the algorithm is explicit computations of complexity measure as treated in [5], [6], and [7]. Examples of both applications can be found in Section 4.

For all the technical details about the iterative scaling algorithm in its traditional version, we refer the interested reader to Chapters 3 and 5 of [1].

All proofs can be found in the Appendix.

1.1 Technical Definitions

We take the same definitions and notations as in [4], except that we let the output be multiple. More precisely, we consider a set of N input nodes V, taking values in the sets $X_1, \ldots, X_N$, and a set of M output nodes W, taking values in the sets $Y_1, \ldots, Y_M$. We write the input globally as $X := X_1 \times \cdots \times X_N$, and the output globally as $Y := Y_1 \times \cdots \times Y_M$. We denote by $F(Y)$ the set of real functions on Y, and by $P(X)$ the set of probability measures on X.

Definition 1. Let $I \subseteq V$ and $J \subseteq W$. We call F_{IJ} the space of functions which only depend on X_I and Y_J:

$$F_{IJ} := \big\{ f \in F(X,Y) \;\big|\; f(x_I, x_{I^c}, y_J, y_{J^c}) = f(x_I, x'_{I^c}, y_J, y'_{J^c})\ \forall x_{I^c}, x'_{I^c}, y_{J^c}, y'_{J^c} \big\} \,. \tag{1}$$

We can model the channel from X to Y as a Markov kernel k, that assigns to each $x \in X$ a probability measure on Y (for a detailed treatment, see [8]). Here we will consider only finite systems, so we can think of a channel simply as a transition matrix (or stochastic matrix), whose rows sum to one.

$$k(x;y) \geq 0 \quad \forall x, y; \qquad \sum_y k(x;y) = 1 \quad \forall x \,. \tag{2}$$

The space of channels from X to Y will be denoted by $K(X;Y)$. We will denote by X and Y also the corresponding random variables, whenever this does not lead to confusion.

Conditional probabilities define channels: if $p(X,Y) \in P(X,Y)$ and the marginal $p(X)$ is strictly positive, then $p(Y|X) \in K(X;Y)$ is a well-defined

channel. Viceversa, if $k \in K(X;Y)$, given $p \in P(X)$ we can form a well-defined joint probability:

$$pk(x,y) := p(x)\, k(x;y) \quad \forall x, y\ . \tag{3}$$

To extend the notion of divergence from probability distributions to channels, we need an "input distribution":

Definition 2. Let $p \in P(X)$, let $k, m \in K(X;Y)$. Then:

$$D_p(k||m) := \sum_{x,y} p(x)\, k(x;y)\, \log \frac{k(x;y)}{m(x;y)}\ . \tag{4}$$

Let p, q be joint probability distributions on $X \times Y$, and let D be the KL-divergence. Then the following "chain rule" holds:

$$D(p(X,Y)||q(X,Y)) = D(p(X)||q(X)) + D_{p(X)}(p(Y|X)||q(Y|X))\ . \tag{5}$$

2 Families of Channels

Suppose we have a family $\mathcal{E}$ of channels, and a channel k that may not be in $\mathcal{E}$. Then we can define the "distance" between k and $\mathcal{E}$ in terms of D_p.

Definition 3. Let p be an input distribution. The divergence between a channel k and a family of channels $\mathcal{E}$ is given by:

$$D_p(k||\,\mathcal{E}) := \inf_{m \in \mathcal{E}} D_p(k||m)\ . \tag{6}$$

If the minimum is uniquely realized, we call the channel

$$\pi_{\mathcal{E}} k := \arg\min_{m \in \mathcal{E}} D_p(k||m) \tag{7}$$

the *rI-projection* of k on $\mathcal{E}$ (and simply "an" rI-projection if it is not unique).

The families considered here are of two types, dual to each other: linear and exponential. For both cases, we take the closures, so that the minima defined above always exist.

Definition 4. A *mixture family* of $K(X;Y)$ is a subset of $K(X;Y)$ defined by one or several affine equations, i.e., the locus of the k which satisfy a (finite) system of equations in the form:

$$\sum_{x,y} k(x;y) f_i(x,y) = c_i\ , \tag{8}$$

for some functions $f_i \in F(X,Y)$, and some constants c_i.

Example. Consider a channel $m \in K(X; Y_1, Y_2)$. We can form the marginal:

$$m(x; y_1) := \sum_{y_2} m(x; y_1, y_2) \ . \tag{9}$$

The channels $k \in K(X; Y_1, Y_2)$ such that $k(x; y_1) = m(x; y_1)$ form a mixture family, defined by the system of equations (for all $x' \in X$, $y_1' \in Y_1$):

$$\sum_{x, y_1, y_2} k(x; y_1, y_2)\, \delta(x, x') \delta(y_1, y_1') = m(x'; y_1') \ , \tag{10}$$

where the function $\delta(z, z')$ is equal to 1 for $z = z'$, and zero for any other case.

More in general, let $\mathcal{L}$ be a (finite-dimensional) linear subspace of $F(X, Y)$, and let $k \in K(X; Y)$. Then:

$$\mathcal{M}(k, \mathcal{L}) := \left\{ m \in K(X; Y) \,\middle|\, \sum_{x,y} m(x; y) l(x, y) = \sum_{x,y} k(x; y) l(x, y) \ \forall l \in \mathcal{L} \right\} \tag{11}$$

is a mixture family, which we call *generated by k and $\mathcal{L}$*.

Definition 5. A (closed) *exponential family* of $K(X; Y)$ is (the closure of) a subset of $K(X; Y)$ of channels in the form:

$$\frac{e^{f(x,y)}}{Z(x)}\, k(x; y) \ , \tag{12}$$

where f satisfies affine constraints, k is fixed, and:

$$Z(x) := \sum_y e^{f(x,y)}\, k(x; y) \tag{13}$$

so that the channel is correctly normalized.

This is a sort of multiplicative equivalent of mixture families, as the exponent satisfies constraints similar to (8).

Example. Let $\mathcal{L}$ be a (finite-dimensional) linear subspace of $F(X, Y)$, and let $k \in K(X; Y)$. Then the closure:

$$\mathcal{E}(k, \mathcal{L}) := \left\{ \frac{e^{l(x,y)}}{Z(x)}\, k(x; y) \,\middle|\, Z(x) = \sum_y e^{l(x,y)}\, k(x; y),\ l \in \mathcal{L} \right\} \tag{14}$$

is an exponential family, which again we call *generated by k and $\mathcal{L}$*.

This family is in some sense "dual" to the family in (11). The duality is expressed more precisely by the following result.

Theorem 1. *Let $\mathcal{L}$ be a subspace of $F(X, Y)$. Let $p \in P(X)$ be strictly positive. Let $k_0 \in K(X; Y)$ be a strictly positive "reference" channel. Let $\mathcal{E} := \mathcal{E}(k_0, \mathcal{L})$ and $\mathcal{M} := \mathcal{M}(k, \mathcal{L})$. For $k' \in K(X; Y)$, the following conditions are equivalent:*

1. $k' \in \mathcal{M} \cap \mathcal{E}$.

2. $k' \in \mathcal{E}$, *and* $D_p(k||k') = \inf_{m \in \mathcal{E}} D_p(k||m)$.

3. $k' \in \mathcal{M}$, *and* $D_p(k'||k_0) = \inf_{m \in \mathcal{M}} D_p(m||k_0)$.

In particular, k' *is unique, and it is exactly* $\pi_{\mathcal{E}} k$.

Geometrically, we are saying that $k' = \pi_{\mathcal{E}} k$, the rI-projection of k on $\mathcal{E}$. We call the mapping $k \to k'$ the *rI-projection operator*, and the mapping $k_0 \to k'$ the *I-projection operator* These are the channel equivalent of the I-projections introduced in [9] and generalized in [10]. The result is illustrated in Figure 1.

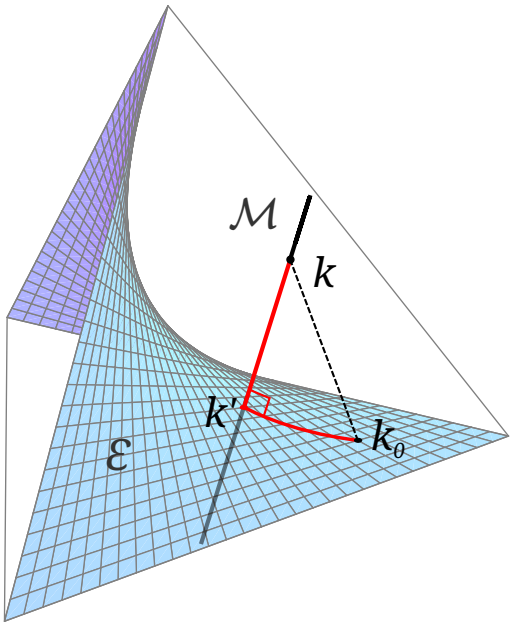

Figure 1: Illustration of Theorem 1. The point k' at the intersection minimizes on $\mathcal{E}$ the distance from k, and minimizes on $\mathcal{M}$ the distance from k_0.

As suggested by Figure 1, I- and rI-projections on exponential families satisfy a Pythagoras-type equality. For any $m \in \mathcal{E}$, with $\mathcal{E}$ exponential family:

$$D_p(k||m) = D_p(k||\pi_{\mathcal{E}} k) + D_p(\pi_{\mathcal{E}} k||m) \ . \tag{15}$$

This statement follows directly from the analogous statement for probability distribution found in [11], after applying the chain rule (5).

3 Algorithm

The algorithm can be considered as a channel equivalent of the iterative scaling procedure for joint distributions, which can be found in Chapter 5 of [1].

Translated into our language, that theorem says the following:

Theorem 2 (Theorem 5.1 of [1]). *Let $\mathcal{L}_1, \ldots, \mathcal{L}_n$ be mixture families of joint distributions with nonempty intersection $\mathcal{L}$. Denote by $\Sigma_i q$ the I-projection of a joint q onto the family $\mathcal{L}_i$. Consider the sequence that starts at q^0 and is defined iteratively by:*

$$q^j := \Sigma_{(j \bmod n)} q^{j-1} \ . \tag{16}$$

Then $\{q^j\}$ converges, and the limit point is the I-projection of q^0 onto $\mathcal{L}$, i.e. if we call:

$$\lim_{i \to \infty} q^i := q \ , \tag{17}$$

then $q \in \mathcal{L}$, and for any $\bar{q} \in \mathcal{L}$:

$$D(q^0||\bar{q}) = D(q^0||q) + D(q||\bar{q}) \ . \tag{18}$$

Our result depends on the theorem above, in the following way. We define a marginal procedure for channels, which in general depends on the choice of an input distribution. We define mixture families of channels with fixed marginals in a way compatible with the equivalent for joints. We then define scalings of channels, and prove that they give the desired result at the joint level. This makes it possible to translate the statement of Theorem 2 to an analogous statement for channels, Theorem 7.

Unless otherwise stated, all the input distributions here will be assumed strictly positive. All our proofs can be found in the appendix.

Definition 6. Consider an input distribution $p \in P(X)$. Let $I \subseteq [N], J \subseteq [M], J \neq \emptyset$. We define the *marginal operator* for channels as:

$$k(x; y) \mapsto k(x_I; y_J) := \sum_{x_{I^c}, y_{J^c}} p(x_{I^c}|x_I)\, k(x_I, x_{I^c}; y_J, y_{J^c}) \ , \tag{19}$$

given the input p.

Proposition 3. *Defined as above, $k(X_I; Y_J)$ is exactly the conditional probability for the marginal $pk(X_I, Y_J)$. In other words, $k(X, Y)$ has marginal $k(X_I; Y_J)$ if and only if $pk(X, Y)$ has marginal $pk(X_I, Y_J)$.*

Definition 7. Consider an input distribution $p \in P(X)$. Let $I \subseteq [N], J \subseteq [M]$, $J \neq \emptyset$. We define the mixture families $\mathcal{M}_{IJ}(\bar{k})$ as:

$$\mathcal{M}_{IJ}(\bar{k}) := \left\{ k(x_{1\ldots n}; y_{1\ldots m}) \,\middle|\, p(x_I)\, k(x_I; y_J) = p(x_I)\, \bar{k}(x_I; y_J) \right\} \ , \tag{20}$$

where the $\bar{k}(x_I; y_J)$ are prescribed channel marginals.

Analogously, let $\bar{q}$ be a probability distribution in $P(X_I, Y_J)$. We define the mixture families:

$$\mathcal{J}_{IJ}(\bar{q}) := \left\{ q(x_{1\ldots n}, y_{1\ldots m}) \,\middle|\, q(x_I, y_J) = \bar{q}(x_I, y_J) \right\} \ . \tag{21}$$

Proposition 3 says that, for any $k \in K(X;Y)$, for any (strictly positive) $p \in P(X)$, and for any $I \subseteq [N], J \subseteq [M]$:

$$k \in \mathcal{M}_{IJ}(\bar{k}) \quad \Leftrightarrow \quad pk \in \mathcal{J}_{IJ}(p\bar{k}) \; . \tag{22}$$

Proposition 4. *$\mathcal{M}_{IJ}(\bar{k})$ is exactly the set $\mathcal{M}(\bar{k}, \mathcal{L})$ of equation (11), where as $\mathcal{L}$ we take the space F_{IJ} of functions which only depend on the nodes in I, J.*

Just as in [1], the I-projections for single marginals can be obtained by scaling. For joint distributions the scaling is done in this way: if $\bar{p}(X_I, Y_J)$ is a "prescribed" marginal, then:

$$\sigma^{\bar{q}}_{IJ}\, p\,(X,Y) := p(X,Y)\, \frac{\bar{q}(X_I, Y_J)}{p(X_I, Y_J)} \tag{23}$$

will have the prescribed marginals, and even be the I-projection of p on $\mathcal{J}_{IJ}(\bar{q})$. i.e., $\sigma^{\bar{q}}_{IJ}\, p \in \mathcal{J}_{IJ}(\bar{q})$, and:

$$D(p||\bar{q}) = D(p||\, \sigma^{\bar{q}}_{IJ}\, p) + D(\sigma^{\bar{q}}_{IJ}\, p||\bar{q}) \; . \tag{24}$$

For the proof, see Chapter 3 and Section 5.1 of [1].

In the case of channels, the scaling is instead done in two steps.

Definition 8. We define the (unnormalized) *IJ-scaling* as the operator $\sigma^{\bar{k}}_{IJ} : K(X;Y) \to F(X,Y)$, mapping k to:

$$\sigma^{\bar{k}}_{IJ}\, k\,(x,y) := k(x_I, x_{I^c}; y_J, y_{J^c}) \frac{\bar{k}(x_I; y_J)}{k(x_I; y_J)} \; . \tag{25}$$

We have that $\sigma^{\bar{k}}_{IJ}\, k$ is *not* an element of $\mathcal{M}_{IJ}$, as in general it is not even in $K(X;Y)$ (i.e. a correctly normalized channel). However, at the joint level this corresponds exactly to the joint scaling:

Proposition 5. *Let $p \in P(X)$, $k \in K(X;Y)$, and $\bar{k} \in K(X_I;Y_J)$. Then:*

$$p(\sigma^{\bar{k}}_{IJ}\, k) = \sigma^{p\bar{k}}_{IJ}\, pk \; . \tag{26}$$

This implies that $p\, \sigma^{\bar{k}}_{IJ}\, k$ is the I-projection of pk to the family $\mathcal{J}_{IJ}(p\bar{k})$.

Definition 9. We define the *normalized IJ-scaling* as the operator $N\, \sigma^{\bar{k}}_{IJ} : K(X,Y) \to K(X;Y)$, mapping k to:

$$N\, \sigma^{\bar{k}}_{IJ}\, k\,(x,y) := \frac{1}{Z(x)}\; \sigma^{\bar{k}}_{IJ}\, k\,(x,y) \; , \tag{27}$$

where:

$$Z(x) := \sum_{y'} \sigma^{\bar{k}}_{IJ}\, k\,(x,y') \; . \tag{28}$$

At the joint level, this corresponds to scaling of the input in the following way:

Proposition 6. *Let $p \in P(X)$, $k \in K(X;Y)$, and $\bar{k} \in K(X_I;Y_J)$. Then:*

$$p(N\,\sigma_{IJ}^{\bar{k}}\,k) = \sigma_{[N]}^{p}\,\sigma_{IJ}^{p\bar{k}}\,pk\ . \tag{29}$$

This implies that $pN\,\sigma_{IJ}^{\bar{k}}\,k$ is the I-projection of $p\,\sigma_{IJ}^{\bar{k}}\,k$ to the family with prescribed input $p(X)$. For brevity, let's call this family $\mathcal{J}_{[N]}(p)$.

Now $N\,\sigma_{IJ}^{\bar{k}}\,k$ is an element of $K(X;Y)$, but still *not* of $\mathcal{M}_{IJ}$. However, if we iterate the operator $N\,\sigma_{IJ}^{\bar{k}}$, the resulting sequence will converge to the projector on $\mathcal{M}_{IJ}$. More in general, the following result holds:

Theorem 7 (Main result)**.** *For $1 \leq i \leq n$, let $I_i \subseteq [N]$ be subsets of $[N]$ and $J_i \subseteq [M]$ be nonempty subsets of $[M]$. Take an input distribution $p \in P(X)$ and a channel $\bar{k} \in K(X,Y)$. Define the mixture families of prescribed marginals:*

$$\mathcal{M}_i = \mathcal{M}_{I_iJ_i}(\bar{k}(X_{I_i},Y_{J_i}))\ , \tag{30}$$

and their intersection, which is also a mixture family (nonempty, as it contains at least $\bar{k}$):

$$\mathcal{M} := \bigcap_i \mathcal{M}_i\ . \tag{31}$$

Choose a (different) channel $k^0 \in K(X,Y)$ and consider the sequence of normalized scalings starting at k^0 and defined iteratively by:

$$k^j := N\sigma_{I_{(j \bmod n)}J_{(j \bmod n)}}k^{j-1}\ . \tag{32}$$

Then:

- *k^i converges to a limit channel:*

$$\lim_{i\to\infty} k^i := l\ ; \tag{33}$$

- *The limit l is the I-projection of k^0 on $\mathcal{M}$, i.e. $l \in \mathcal{M}$ and:*

$$D_p(k^0||\bar{k}) = D_p(k^0||l) + D_p(l||\bar{k})\ . \tag{34}$$

The proof can be found in the appendix.

To apply the Theorem 7 in our algorithm, we choose as initial channel k^0 exactly the reference channel k_0 of Theorem 1, usually the uniform channel. As $\bar{k}$ we take exactly the "prescription channel" k of Theorem 1, i.e. the channel which has the desired marginals. The result of the iterative scaling will be the rI-projection of k on the desired exponential family.

4 Applications

4.1 Synergy Measures

The algorithm presented here permits to compute the decompositions of mutual information between inputs and outputs in [3] and [4]. We give here examples of

computations of *pairwise synergy* as an rI-projection for channels, as described in [4]. It is not within the scope of this article to motivate this measure, we rather want to show how it can be computed.

Let k be a channel from $X = (X_1, X_2)$ to Y. Let $p \in P(X)$ be a strictly positive input distribution. We define in [4] the synergy of k as:

$$d_2(k) := D_p(k||\,\mathcal{E}_1)\;, \tag{35}$$

where $\mathcal{E}_1$ is the (closure of the) family of channels in the form:

$$m(x_1, x_2; y) = \frac{1}{Z(X)} \exp\big(\phi_0(x_1, x_2) + \phi_1(x_1, y) + \phi_2(x_2, y)\big)\;, \tag{36}$$

where:

$$Z(x) := \sum_y \exp\big(\phi_0(x_1, x_2) + \phi_1(x_1, y) + \phi_2(x_2, y)\big)\;, \tag{37}$$

and:

$$\phi_0 \in F_{\{1,2\}\emptyset}\,, \quad \phi_1 \in F_{\{1\}\{1\}}\,, \quad \phi_2 \in F_{\{2\}\{1\}}\;. \tag{38}$$

According to Theorem 1, the rI-projection of k on $\mathcal{E}_1$ is the unique point k' of $\mathcal{E}_1$ which has all the prescribed marginals:

$$k'(x_1; y) = k(x_1; y)\,, \quad k'(x_2; y) = k(x_2; y)\;, \tag{39}$$

and can therefore be computed by iterative scaling, either of the joint distribution (as it is traditionally done, see [1]), or of the channels (our algorithm).

Here we present a comparison of the two algorithms, implemented similarly and in the same language (Mathematica). The red dots represent our (channel) algorithm, and the blue dots represent the joint rescaling algorithm.

For the easiest channels (see Figure 2), both algorithm converge instantly.

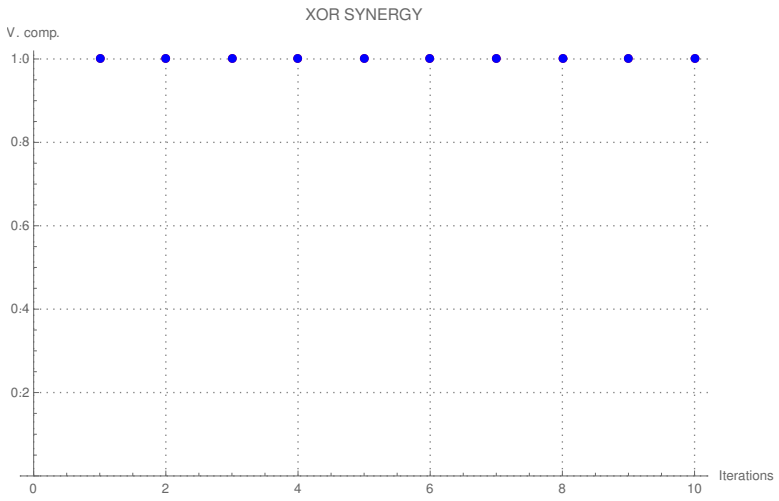

Figure 2: Comparison of convergence times for the synergy of the XOR gate. Both algorithms get immediatly the desired result. (The dots here are overlapping, the red ones are not visible.)

A more interesting example is a randomly generated channel (Figure 3), in which both method need 5-10 iterations to get to the desired value. However, the channel method is slightly faster.

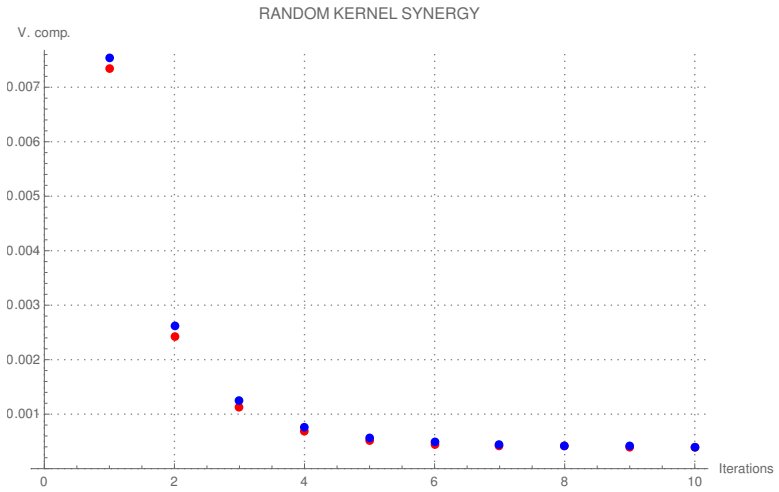

Figure 3: Comparison of convergence times for the synergy of a randomly generated channel. The channel method (red) is slightly faster.

The most interesting example is the synergy of the AND gate, which should be zero according to the procedure [4]. In that article, we mistakenly wrote a

different value, that here we would like to correct (it is zero). The convergence to zero is very slow, of the order of $1/n$ (Figure 4). It is clearly again slightly faster for the channel method in terms of iterations.

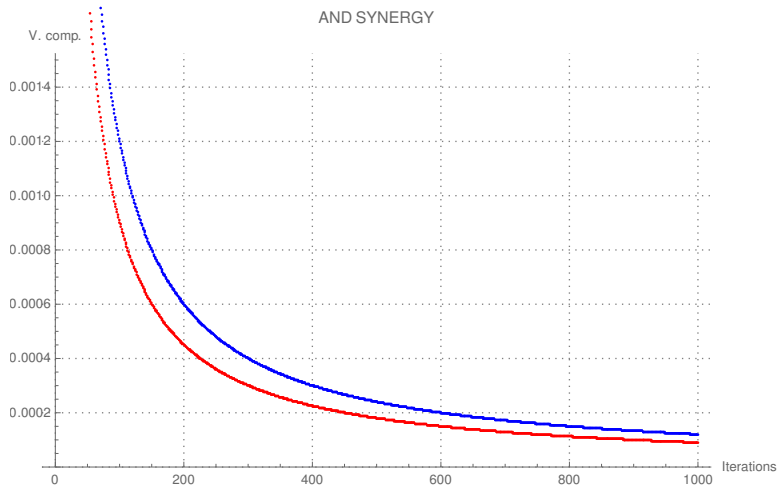

Figure 4: Comparison of convergence times for the synergy of the AND gate. The channel method (red) tends to zero proportionally to $n^{-1.05}$, the joint method (blue) proportionally to $n^{-0.95}$.

It has to be noted, however, that rescaling a channel requires more elementary operations than rescaling a joint distribution. Because of this, one single iteration with our method takes longer than with the joint method. (As explained in Section 3, a scaling for the channel corresponds to two scalings for the joint.) In the end, despite the need of fewer iterations, the total computation time of a projection with our algorithm can be longer (depending on the particular problem). For example, again for the synergy of the AND gate, we can plot the computation time as a function of the accuracy (distance to actual value), down to 10^{-3}. The results are shown in Figure 5.

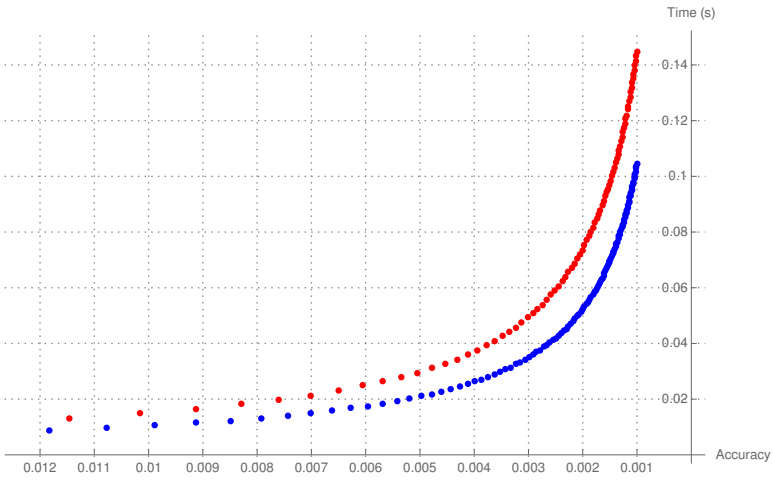

Figure 5: Comparison of total computation times for the synergy of the AND gate. The channel method (red) is slightly slower than the joint method (blue).

To get to the same accuracy, though, the channel approach used less iterations. In summary, our algorithm is better in terms of iteration complexity, but generally worse in terms of computing time.

4.2 Complexity Measures

Iterative scaling can also be used to compute measures of complexity, as defined in [5], [6], and in Section 6.9 of [7]. For simplicity, consider two inputs X_1, X_2, two outputs Y_1, Y_2 and a generic channel between them. In general, any sort of interaction is possible, which in terms of graphical models (see [12]) can be represented by diagrams such as those in Figure 6.

Any line in the graph indicates an interaction between the nodes. In [5] the outputs are assumed to be conditionally independent, i.e. they do not directly interact (or, their interaction can be *explained away* by conditioning on the inputs). In this case the graph looks like Figure 6a, and the maginals to preserve are those of the family of pairs (X_{I_i}, Y_{J_i}), $i = 1, 2$ with: $X_{I_1} = X_{\{1,2\}}$, $Y_{J_1} = Y_{\{1\}}$, $X_{I_2} = X_{\{1,2\}}$, $Y_{J_2} = Y_{\{2\}}$.

Suppose now that Y_1, Y_2 correspond to X_1, X_2 at a later time. In this case it is natural to assume that the system is not complex if Y_1 does not depend (directly) on X_2, and Y_2 does not depend (directly) on X_1. Intuitively, in this case "the whole is exactly the sum of its parts". In terms of graphical models, this means that our system is represented by Figure 6b, meaning that the subsets of nodes in question are now only the ones given by $X_{I_1} = X_{\{1\}}$, $Y_{J_1} = Y_{\{1\}}$, $X_{I_2} = X_{\{2\}}$, $Y_{J_2} = Y_{\{2\}}$. These channels (or joints) form an exponential family (see [5]) which we call $\mathcal{F}_1$.

Suppose now, though, that the outputs are not conditionally independent anymore, because of some "noise" (see [6] and [7]). This way the interaction

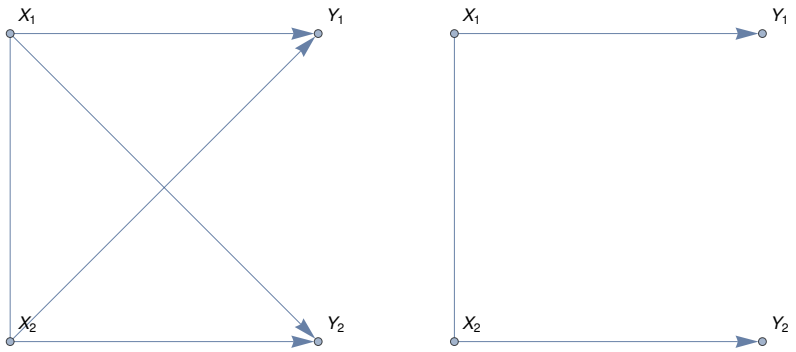

Figure 6: a) The graphical model corresponding to conditionally independent outputs Y_1 and Y_2 are indeed correlated, but only indirectly, via the inputs. b) The graphical model corresponding to a non-complex system.

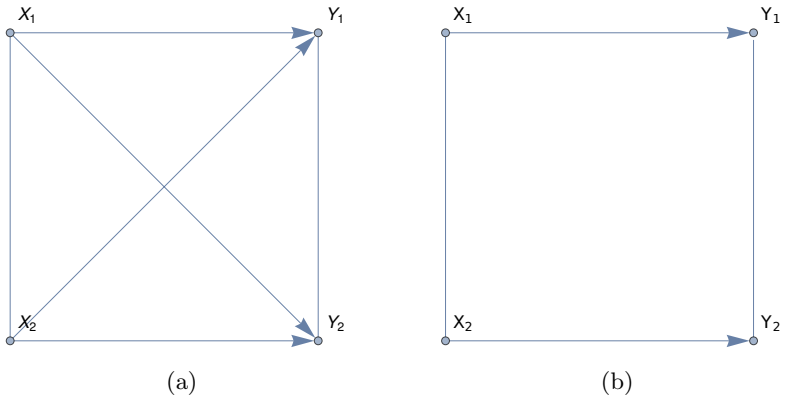

Figure 7: a) The graphical model of Figure 6a, with correlation between the outputs. b) The non-complex model of Figure 6b, with correlation between the outputs.

structure would look like Figure 7, i.e. the "complete" subset given by (X_I, Y_J) with $X_I = X_{\{1,2\}}$ and $Y_J = Y_{\{1,2\}}$. In particular, a non-complex but "noisy" system would be represented by Figure 7b, and have subsets of nodes given by the pairs (X_{I_i}, Y_{J_i}), $i = 1, 2, 3$ with: $X_{I_1} = X_{\{1\}}$, $Y_{J_1} = Y_{\{1\}}$, $X_{I_2} = X_{\{2\}}$, $Y_{J_2} = Y_{\{2\}}$, $X_{I_3} = X_{\emptyset}$, $Y_{J_3} = Y_{\{1,2\}}$. Such channels form again an exponential family, which we call $\mathcal{F}_2$.

We would like now to have a measure of complexity for a channel (or joint). In [5], the measure of complexity is defined as the divergence from the family $\mathcal{F}_1$ represented in Figure 6b. We will call such a measure c_1. In case of noise, however, it is argued in [6] and [7] that the divergence should be computed from the family $\mathcal{F}_2$ represented in 7b (for example, as written in the cited papers, because such a complexity measure should be required to be upper bounded by the mutual information between X and Y). We will call such a measure c_2.

Both divergences can be computed with our algorithm. As an example, we have considered the following channel:

$$k(x_1, x_2, x_3; y_1, y_2) = \frac{1}{Z(x)} \exp\left(\big(\alpha\, x_1\, x_2 + \beta x_3\big)(y_1 - y_2)\right), \tag{40}$$

with:

$$Z(x) = \sum_{y_1', y_2'} \exp\left(\big(\alpha\, x_1\, x_2 + \beta x_3\big)(y_1' - y_2')\right). \tag{41}$$

Here X_3 represents a node of "unknown input noise" that adds correlation between the outputs (of unknown form) when if it is not observed. We have chosen $\alpha = 1$ and $\beta = 2$, and a uniform input probability p. After marginalizing out X_3 (obtaining then an element of the type of Figure 7a), we can compute the two divergences:

- $c_1(k) = D_p(k||\,\mathcal{F}_1) = 0.519$.
- $c_2(k) = D_p(k||\,\mathcal{F}_2) = 0.110$.

This could indicate that c_1 is incorporating part the correlation of the output nodes due to the "noise", and therefore probably overestimating the complexity, at least in this case.

One could nevertheless also argue that c_2 can underestimate complexity, as we can see in the following "control" example. Consider the channel:

$$h(x_1, x_2; y_1, y_2) = \frac{1}{Z(x)} \exp\left(\big(\alpha\, x_1\, x_2\big)(y_1 - y_2)\right), \tag{42}$$

with:

$$Z(x) = \sum_{y_1', y_2'} \exp\left(\big(\alpha\, x_1\, x_2\big)(y_1' - y_2')\right), \tag{43}$$

which is represented by the graph in Figure 6a. If the difference between c_1 and c_2 were just due to the noise, then for our new channel $c_1(h)$ and $c_2(h)$ should be equal. This is not the case:

- $c_1(h) = D_p(h||\mathcal{F}_1) = 0.946$.
- $c_1(h) = D_p(h||\mathcal{F}_2) = 0.687$.

The divergences are still different. This means that there is an element h_2 in $\mathcal{F}_2$, which does *not* lie in $\mathcal{F}_1$, for which:

$$D_p(h||h_2) < D_p(h||h_1) \quad \forall h_1 \in f_1 \ . \tag{44}$$

The difference is this time smaller, which could mean that noise still does play a role, but in rigor it is hard to say, since none of these quantities is linear, and divergences do not satisfy a triangular inequality.

We do not want to argue here in favor or against any of these measures. We would rather like to point out that such considerations can be done mostly after explicit computations, which can be done with iterative scaling.

References

[1] Csiszár, I. and Shields, P. C. *Information Theory and Statistics: A Tutorial*. Foundations and Trends in Communications and Information Theory, 1(4):417–528, 2004

[2] Goodman, J. *Sequential Conditional Generalized Iterative Scaling*. Proceedings of the 40th Annual Meeting of ACL:9–16, 2002.

[3] Olbrich, E., Bertschinger, N., and Rauh, J. *Information decomposition and synergy*. Entropy, 17(5):3501–3517, 2015

[4] Perrone, P. and Ay, N. *Hierarchical quantification of synergy in channels*. Frontiers in Robotics and AI, 35(2), 2016.

[5] Ay, N. *Information Geometry on Complexity and Stochastic Interaction*. Entropy, 17, 2432–2458, 2015.

[6] Oizumi, M., Tsuchiya, N., and Amari, S. *A unified framework for information integration based on information geometry* Preprint available on arXiv:1510.04455, 2015.

[7] Amari, S. *Information Geometry and Its Applications* Springer, 2016.

[8] Kakihara, Y. *Abstract Methods in Information Theory*. World Scientific, 1999.

[9] Csiszár, I. *I-divergence geometry of probability distributions and minimization problems*. Annals of Probability, 3:146–158, 1975.

[10] Csiszár, I. and Matúš, F. *Information projections revisited*. IEEE Transactions on Information Theory, 49:1474–1490, 2003.

[11] Amari, S. *Information geometry on a hierarchy of probability distributions*. IEEE Transactions on Information Theory, 47(5):1701–1709, 2001.

[12] Lauritzen, S. L. *Graphical Models*. Oxford, 1996.

[13] Williams, P. L. and Beer, R. D. *Nonnegative decomposition of multivariate information*. Submitted. Preprint available on arXiv:1004.2151, 2010.

[14] Amari, S. and Nagaoka, H. *Differential geometry of smooth families of probability distributions*. Technical Report METR 82-7, Univ. of Tokyo, 1982.

[15] Amari, S. and Nagaoka, H. *Methods of Information Geometry*. Oxford, 1993.

A Proofs

Proof of Theorem 1. $1 \Leftrightarrow 2$: Choose a basis $f_1, \dots, f_d$ of $\mathcal{L}$. Define the map $\theta \mapsto k_\theta$, with:

$$k_\theta(x;y) = k(\theta_1, \dots, \theta_d)(x;y) := \frac{1}{Z_\theta(x)}\, k_0(x;y)\, \exp\left(\sum_{j=1}^{d} \theta_j f_j(x,y)\right) \;, \quad (45)$$

and:

$$Z_\theta(x) := \sum_y k_0(x;y)\, \exp\left(\sum_{i=1}^{d} \theta_i f_i(x,y)\right) \;. \quad (46)$$

Then:

$$D_p(k||k_\theta) = D_p(k||k_0) - \sum_{j=1}^{d} \theta_j\; \mathbb{E}_{pk}[f_j] + \mathbb{E}_p[\log Z_\theta] \;. \quad (47)$$

Deriving (where ∂_j is w.r.t. θ_j):

$$\partial_j\, D_p(k||k_\theta) = -\,\mathbb{E}_{pk}[f_j] + \mathbb{E}_p\left[\frac{\partial_j\, Z_\theta}{Z_\theta}\right] \;. \quad (48)$$

The term in the last brackets is equal to:

$$\frac{\partial_j\, Z_\theta}{Z_\theta} = \frac{1}{Z_\theta} \sum_y k_0(x;y)\, \exp\left(\sum_{i=1}^{d} \theta_i f_i(x,y)\right) f_j(x,y) \quad (49)$$

$$= \sum_y k_\theta(x;y) f_j(x,y) \;, \quad (50)$$

so that (48) now reads:

$$\partial_j\, D_p(k||k_\theta) = -\,\mathbb{E}_{pk}[f_j] + \mathbb{E}_{pk_\theta}[f_j] \;. \quad (51)$$

This quantity is equal to zero for every j if and only if $k_\theta \in \mathcal{M}$. Now if k_θ is a minimizer, it satisfies (51), and so $k_\theta \in \mathcal{M}$. Viceversa, suppose $k_\theta \in \mathcal{M}$, so that it satisfies (51) for every j. To prove that it is a global minimizer, we look at the Hessian:

$$\partial_i\, \partial_j\, D_p(k||k_\theta) = \partial_i\, \partial_j\, D(pk||pk_\theta) \;. \quad (52)$$

This is precisely the covariance matrix of the joint probability measure pk_θ, which is positive definite.

$1 \Leftrightarrow 3$: For every $m \in \mathcal{M}$, we have:

$$D_p(m||k_0) = \sum_{x,y} p(x)\, m(x;y) \log \frac{m(x;y)}{k_0(x;y)} = \mathbb{E}_{pm}\left[\log \frac{m}{k_0}\right] \;. \quad (53)$$

If $k' \in \mathcal{E}$, then:

$$D_p(m||k_0) = \mathbb{E}_{pm}\left[\log \frac{m}{k'} + \log \frac{k'}{k_0}\right] = D_p(m||k') + \mathbb{E}_{pm}\left[\log \frac{k'}{k_0}\right] \;. \quad (54)$$

By definition of $\mathcal{E}$, the logarithm in the last brackets belongs to $\mathcal{L}$, and since $m \in \mathcal{M}$:

$$\mathbb{E}_{pm}\left[\log \frac{k'}{k_0}\right] = \mathbb{E}_{pk}\left[\log \frac{k'}{k_0}\right] = \mathbb{E}_{pk'}\left[\log \frac{k'}{k_0}\right] . \tag{55}$$

Inserting in (54):

$$D_p(m||k_0) = D_p(m||k') + \mathbb{E}_{pk'}\left[\log \frac{k'}{k_0}\right] = D_p(m||k') + D_p(k'||k_0) . \tag{56}$$

Since $D_p(m||k') \geq 0$, (56) shows that k' is a minimizer. Since $D_p(m||k_0) = D(pm||pk_0)$ is strictly convex in the first argument, its minimizer is unique. □

Proof of Proposition 3.

$$\begin{aligned} p(x_I)k(x_I; y_J) &= p(x_I) \sum_{x_{I^c}, y_{J^c}} p(x_{I^c}|x_I)\, k(x_I, x_{I^c}; y_J, y_{J^c}) && (57)\\ &= \sum_{x_{I^c}, y_{J^c}} p(x_I)\, p(x_{I^c}|x_I)\, k(x_I, x_{I^c}; y_J, y_{J^c}) && (58)\\ &= \sum_{x_{I^c}, y_{J^c}} p(x_I, x_{I^c})\, k(x_I, x_{I^c}; y_J, y_{J^c}) && (59)\\ &= \sum_{x_{I^c}, y_{J^c}} pk(x_I, x_{I^c}, y_J, y_{J^c}) && (60)\\ &= pk(x_I, y_J) . && (61) \end{aligned}$$

□

Proof of Proposition 4. For f in F_{IJ}:

$$\mathbb{E}_{pk}[f] = \sum_{x,y} p(x)\, k(x; y)\, f(x, y) = \sum_{x_I, y_J} p(x_I)\, k(x_I; y_J)\, f(x_I; y_J) , \tag{62}$$

and just as well:

$$\mathbb{E}_{p\bar{k}}[f] = \sum_{x,y} p(x)\, \bar{k}(x; y)\, f(x, y) = \sum_{x_I, y_J} p(x_I)\, \bar{k}(x_I; y_J)\, f(x_I; y_J) . \tag{63}$$

The definition in (11) (with strict positivity of p) requires exactly that:

$$\mathbb{E}_{pk}[f] = \mathbb{E}_{p\bar{k}}[f] \tag{64}$$

for every $f \in F_{IJ}$. Using (62) and (63), the equality becomes:

$$\sum_{x_I, y_J} p(x_I)\, k(x_I; y_J)\, f(x_I; y_J) = \sum_{x_I, y_J} p(x_I)\, \bar{k}(x_I; y_J)\, f(x_I; y_J) \tag{65}$$

for every f in F_{IJ}, which means that $k(x_I; y_J) = \bar{k}(x_I; y_J)$. □

Proof of Proposition 5. We have:

$$p\,\sigma^{\bar{k}}_{IJ}\,k\,(x_I,y_J) = p\,\sigma^{\bar{k}}_{IJ}\,k\,(x_I,x_{I^c},y_J,y_{J^c}) \tag{66}$$

$$= p(x_I,x_{I^c})k(x_I,x_{I^c};y_J,y_{J^c})\,\frac{\bar{k}(x_I;y_J)}{k(x_I;y_J)} \tag{67}$$

$$= pk(x_I,x_{I^c},y_J,y_{J^c})\,\frac{p(x_I)\,\bar{k}(x_I;y_J)}{p(x_I)\,k(x_I;y_J)} \tag{68}$$

$$= pk(x_I,x_{I^c},y_J,y_{J^c})\,\frac{p\bar{k}(x_I,y_J)}{pk(x_I,y_J)} \tag{69}$$

$$= \sigma^{p\bar{k}}_{IJ}pk\,(x,y)\;. \tag{70}$$

□

Proof of Proposition 6. The first member is equal to:

$$p(x)N\,\sigma^{\bar{k}}_{IJ}\,k(x,y) = p(x)\,\frac{\sigma_{IJ}\,k\,(x,y)}{\sum_{y'}\sigma^{\bar{k}}_{IJ}\,k\,(x,y')} \tag{71}$$

$$= p(x)\,\frac{\sigma^{p\bar{k}}_{IJ}\,pk(x)\;\sigma_{IJ}\,k\,(x,y)}{\sum_{y'}\sigma^{p\bar{k}}_{IJ}\,pk(x)\;\sigma^{\bar{k}}_{IJ}\,k\,(x,y')} \tag{72}$$

$$= p(x)\,\frac{\sigma^{p\bar{k}}_{IJ}\,pk\,(x,y)}{\sum_{y'}\sigma^{p\bar{k}}_{IJ}\,pk\,(x,y')} \tag{73}$$

$$= \sigma^{p\bar{k}}_{IJ}\,pk\,(x,y)\,\frac{p(x)}{\sigma^{p\bar{k}}_{IJ}\,pk\,(x)} \tag{74}$$

$$= \sigma^{p}_{[N]}\,\sigma^{p\bar{k}}_{IJ}\,pk\;. \tag{75}$$

□

Proof of Theorem 7. In the hypothesis and notation of Theorem 2, take the collection $\mathcal{L}_1,\ldots,\mathcal{L}_{2n}$ in the following way, for $i=1,\ldots,n$:

- $\mathcal{L}_{2i-1} := \mathcal{J}_{[N]}(p)$;
- $\mathcal{L}_{2i} := \mathcal{J}_{I_iJ_i}(p\bar{k})$.

Their intersection $\mathcal{L}$ is nonempty, as it contains at least $p\bar{k}$.

Take as initial distribution $q^0 := pk^0$ and form as in Theorem 2 the sequence $\{q^j\}$ of I-projections. According to Theorem 2, this sequence converges to the I-projection of pk^0 on $\mathcal{L}$. Since $\mathcal{L}\subseteq\mathcal{J}_{[N]}(p)$, this projection will have input marginal equal to $p(X)$, and so we can write it as $p(X)\,l(X;Y)$ for some uniquely defined channel l. We have, for $j\to\infty$:

$$q^j \to pl\;, \tag{76}$$

so in particular, for the subsequence of even-numbered terms also:

$$q^{2j} \to pl\;. \tag{77}$$

This subsequence is defined iteratively by:

$$q^{2j} = \sigma^{p}_{[N]}\, \sigma^{p\bar{k}}_{I_j J_J}\, q^{2(j-1)} \ . \tag{78}$$

Propositions 5 and 6 imply then that:

$$q^{2j} = p\, k^j \tag{79}$$

for every j, where $\{k^j\}$ is the sequence defined in the statement of Theorem 7. Therefore this sequence converges:

$$k^j \to l \ . \tag{80}$$

Since $pl \in \mathcal{L} \subseteq \mathcal{L}_i$ for all i, $l \in \mathcal{M}_i$ for all i because of (22), which by definition means that $l \in \mathcal{M}$. Moreover, pl is the I-projection q of pk^0 on $\mathcal{L}$, which means that:

$$D(pk^0||p\bar{k}) = D(pk^0||pl) + D(pl||p\bar{k}) \ . \tag{81}$$

Using the chain rule of the KL-divergence (5), we get:

$$D_p(k^0||\bar{k}) = D_p(k^0||l) + D_p(l||\bar{k}) \ , \tag{82}$$

which means that l is the I-projection of k^0 on $\mathcal{M}$. □